Spring 2002

# Gridpoints

The Quarterly Publication of the NASA Advanced Supercomputing Division

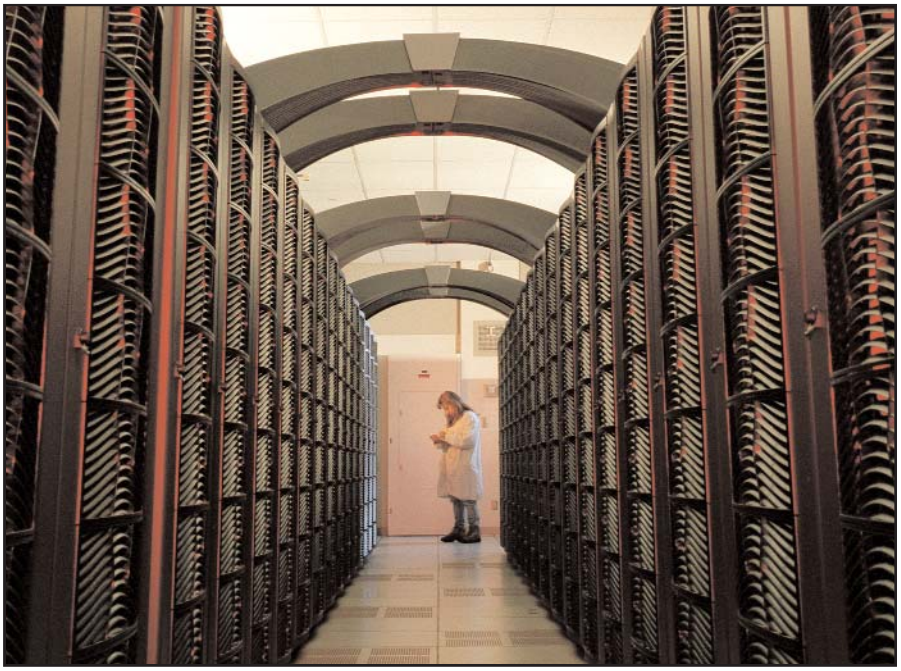

**The evolution of high-performance computers at the NAS Facility is chronicled beginning on page 1A**

FEATURES

# Features

4

## Launching Jobs Into Grid Space

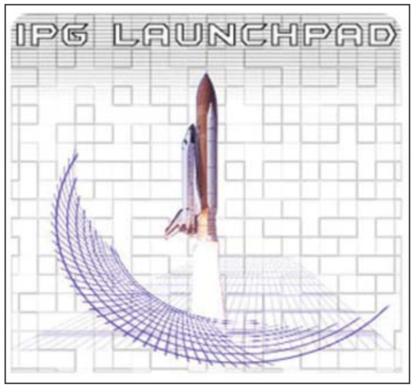


*Information Power Grid users can now send jobs to any IPG resource through the recently developed grid portal, Launch Pad.*
Holly A. Amundson

1A

## The Evolution of High-Performance Computers at the NAS Facility

*The NAS Division has been a leader in high-performance computing since its inception in the mid-1980s. In this special section we examine the supercomputers operated by the division.*
Nicholas A. Veronico

6

## Identifying Organic Molecules: The Link to Our Beginnings

*Computational chemists in the NAS Division are developing a theoretical model to generate accurate data of amino acids, which will help astronomers identify molecules observed in the interstellar medium.*
Holly A. Amundson

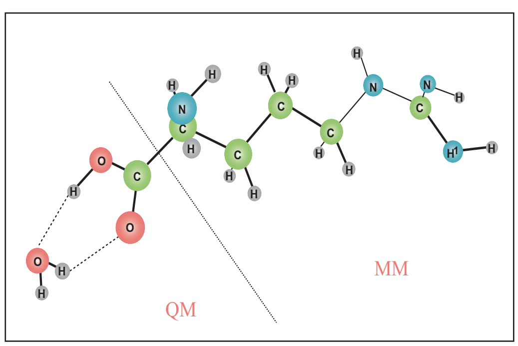


### On The Cover:

*NAS operates a 1,024-processor SGI Origin 3800, the largest single system image high-performance computer in the world. This is the most powerful computer used by NASA scientists to date. For more, turn to page 1A. (NASA/Tom Trower)*

Gridpoints
The Quarterly Publication of the NASA Advanced Supercomputing Division
Vol. 3, No. 1 • Spring 2002

**Executive Editor:** John Ziebarth
**Editor:** Nicholas A. Veronico
**Writers:** Holly A. Amundson, Jill A. Dunbar
**Graphic Designer:** Sue Kim
**Illustrator:** Gail Felchle
**Photographer:** Michael Boswell

**Contact the Editor:**
Nicholas A. Veronico
*Gridpoints*
M/S 258-6
Moffett Field, CA 94035-1000
Tel. (650) 604-3046
Fax (650) 604-4377
nveronico@mail.arc.nasa.gov

The staff of *Gridpoints* would like to thank the following people for their assistance in preparing this publication: Dick Anderson, Bill Arasin, Ron Bailey, Galina Chaban, Bob Ciotti, Jill Dunbar, Chris Gong, T.R. Govindan, John Hardman, Tom Hinke, Winifred Huo, Dennis Jespersen, Julie Jervis, Cetin Kiris, Chris Kleiber, Art Lazanoff, Tony Lisotta, Gina Morello-Ficcadenti, George Myers, Alan Powers, Marcia Redmond, Stuart Rogers, Ryan Spaulding, Ron Sykora, Jim Taft, Leigh Ann Tanner, Mark Tangney, Malu Trehan, Bill Van Dalsem, Rob Van Der Wijngaart, Cliff Williams

*Gridpoints* is a quarterly publication of the NASA Advanced Supercomputing Division at NASA Ames Research Center. The division's staff includes government employees and contractors from Advanced Management Technology Inc., Computer Sciences Corp., Cray Inc., Eloret, Raytheon Co., and SGI.


**Subscribe to *Gridpoints***
Subscriptions to *Gridpoints* are complimentary. Please send e-mail with your name, company name, title, address, and zip code to:
gridpoints@nas.nasa.gov
or subscribe online at:
www.nas.nasa.gov/gridpoints

**Note:** Subscribers outside the continental United States can access *Gridpoints* on the web in PDF format at: http://www.nas.nasa.gov/gridpoints

# News From NAS

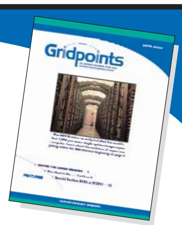


## Boeing Analyzes Sonic Cruiser Using NASA CFD Software

The Boeing Company is now using an array of NASA-developed software packages to analyze different designs of their Sonic Cruiser airplane: the NASA Overset CFD Software (which includes the Overflow flow solver), Chimera Grid Tools, and Pegasus5.

As reported in the February 11, 2002 issue of *Aviation Week and Space Technology* (page 47), Boeing is extensively using computational fluid dynamics (CFD) to analyze the airplane's design, reducing the amount of wind tunnel testing needed.

Boeing plans to use only four high-speed wind tunnel testing sessions during the design of the Sonic Cruiser – two or three fewer sessions than would be required without the use of CFD. Fewer tests translates to significant savings of both time and money – a round of wind tunnel testing can take up to one year to complete.

"This represents significant cost savings to Boeing's multi-billion dollar Sonic Cruiser aircraft design effort. It will also mean that the aircraft will be able to enter service much sooner," said Stuart Rogers, aerospace engineer and co-developer of Pegasus5 and the Chimera

*Continued on page 2*

## NAS Mission

**To lead the country in the research, development, and delivery of revolutionary, high-end computing services and technologies, such as applications and algorithms, tools, system software, and hardware to facilitate NASA mission success.**

# From The Division Chief

*This is my last column as NAS Division Chief as I will be leaving NASA to become Los Alamos National Laboratories' division director of Computer and Computational Sciences. This career change is quite an opportunity for me, but I have very mixed feelings about leaving such a successful organization, and the colleagues and friends with whom I have worked. I'm very proud of the NAS division, its history and legacy, and am honored to have been entrusted with managing the division during the last three and a half years. Let me review some of the things that the NAS staff has accomplished, and the reason that I am proud to have been part of the division.*

*Starting from the early days of computational aeronautical design in the late 1960s and '70s, precursor organizations at Ames Research Center married numerical solutions of the Navier-Stokes equations with the largest computers of the day – the ILLIAC IV, various Control Data Corp. machines, and cousins from IBM and others. This was really the birthplace of numerical fluid dynamics, and laid the groundwork for the next 20 years. The division was formed in the early 1980s from these efforts and introduced the era of vector supercomputing to computational fluid dynamics (CFD). Many of the well-known procedures and services for running a large scientific supercomputing center were developed right here and presented to the rest of the world as a by-product of pushing the limits in computing, while machines like the entire family of Crays graced our shop floor.*

*This "pathfinding" role took NAS into the era of parallel supercomputing starting in the late 1980s and early '90s. Researchers explored a variety of paths to determine the future of scientific computing, experimenting with machines from Intel, SGI, the Thinking Machines Corp. and IBM (see the special section on NAS computers beginning on page 1A). The most promising research computers were then made available to thousands of scientists and engineers across the country.*

*In the last several years, the science and engineering performed on these machines has broadened far beyond CFD. The expertise gained from CFD has been applied in many computational sciences and multiplied the value of the original aeronautics work many-fold. Eminent Earth scientists are being assisted in global climate modeling; computational biologists are exploring the origins of life; and nanotechnology researchers are investigating the construction of devices in this new area of science. All of this is made possible by developments in the science of high-performance computing – ranging from new hardware architectures, middleware, visualization and data tools, and new algorithms and applications.*

*The NAS Division continues its quest for new developments in computer science and is forging the way in two new areas that will be central to the next decade: shared-memory parallel supercomputers and distributed high-performance computing. Both of these areas complement each other as we are now seeing the largest, most powerful supercomputers drawn together in a greater distributed scientific resource – the grid.*

*As I take my leave from NAS, I want everyone to know that I am proud to have been part of this organization, and I see it continuing to be a leading player in the field of high-performance computing.*

*Bill Feiereisen*

*Continued from page 1*

Grid Tools software, from the NASA Advanced Supercomputing (NAS) Division.

NASA Overset CFD software has been used to run hundreds of CFD calculations in the past year at Boeing. The high quality of the software, coupled with the automated CFD processes, makes it possible to turn around design test results quickly and accurately. "We had a lot of confidence in the CFD, and the first high-speed wind tunnel test results were very close to what CFD predicted," explained Walt Gillette, Boeing Sonic Cruiser program manager.

## NAS Scientist Co-develops Unique Load-Balancing Solution

NAS Division scientist Rupak Biswas has co-developed a novel load-balancing algorithm for scientific applications. The method is described in "Parallel Processing of Adaptive Meshes with Load Balancing," which appeared in the December 2001 issue of *IEEE Transactions on Parallel and Distributed Systems* (TPDS). Biswas collaborated with Sajal Das and Daniel Harvey, University of Texas at Arlington. The development team's goal was to create a load-balancing algorithm that also minimized interprocessor communication and data redistribution delays for unstructured grids.

Many scientific applications involving grids lack a uniform underlying structure. These applications are often also dynamic in nature in that the grid structure changes significantly between successive phases of execution on computers.

Mesh adaptation of unstructured grids through selective refinement and coarsening has proven effective in solving complex computational problems. However, an approach using parallel processors is complicated because of frequent load-balancing requirements. Traditional dynamic load balancers are for the most part inadequate because they lack a global view of system loads across processors. The team's paper describes a general-purpose load balancer that utilizes symmetric broadcast networks (SBN) as the underlying communication topology.

Comparisons with PLUM (Parallel Load-balancing for adaptive Unstructured Meshes), a NAS-developed global load-balancing framework, shows that this SBN-based load balancer achieves lower redistribution overhead for dynamic irregular applications by overlapping processing and data migration.

## Cart3D Used to Study Man-Portable Surface-to-Air Missiles

Scientists in the NAS Division's Applications Branch have extended the capabilities of NASA's Cart3D aerodynamic analysis software to simulate the unsteady flight of infrared-guided, man-portable surface-to-air missiles (SAMs). The simulations are being used to analyze the aerodynamic performance characteristics of these SAMs.

In the late 1970s, these portable SAMs proved highly effective against the Soviets in Afghanistan. Since then, these weapons have been sold by the thousands on the black market to various terrorist organizations around the world. While they are very inexpensive to purchase, these missiles have unique flight characteristics that make them difficult to simulate computationally.

NAS researchers Michael Aftosmis and Scott Murman conducted steady and unsteady simulations of controlled flight for representative SAMs on both the 1,024-processor and 512-CPU SGI Origin 3000 systems at the NAS facility at Ames Research Center. More than 720 steady simulations were run on meshes totaling almost 2.4 billion cells.

Two years ago, the Defense Intelligence Agency (DIA) Missile and Space Intelligence Center enlisted the help of NASA and other research organizations to numerically simulate such missiles. The information gained will enable DIA to develop effective measures for countering the threat to civilian and military aircraft posed by terrorists armed with such missiles.

"Results from preliminary NAS studies presented to the DIA have been met with enthusiastic support," said Aftosmis. The steady and unsteady results have shed new light on the flight characteristics of these missiles. Analysts at DIA have relied on early versions of NASA's Cart3D package for several years, and they have already begun to incorporate results from the preliminary NAS Division studies with their in-house studies. Analyses like these allow DIA personnel to plan effective evasive maneuvers for aircraft targeted by these missiles, and support other counter-terrorist work.

The Cart3D package permits end-to-end simulation of inviscid three-dimensional flows around complex vehicles. The parallel analysis code in Cart3D ("flowCart") was developed under NASA's former High Performance Computing and Communications Program, and is well suited to these computations because it scales very well on parallel shared-memory hardware. In addition, many computations can be launched simultaneously to rapidly fill out test matrices. Cart3D was developed at Ames and New York University and is available through the NASA Ames Commercial Technology Office.

## IPG Workshop Highlights Collaboration, Web Services

Researchers from across the U.S. gathered recently to discuss the current state of NASA's geographically distributed computational network, the Information Power Grid (IPG). At the third annual IPG workshop, held

December 4-5 in Palo Alto, Calif., more than 100 computer scientists from academic institutions, government labs, and NASA centers learned about recent developments in grid technology. Thirty scientists presented the state of their research and the development progress of IPG infrastructure components. Peer feedback, problem solving, and networking allowed researchers to address problems encountered in developing new web tools and applications for the IPG.

Kicking off the two-day event, former NAS Division Chief Bill Feiereisen emphasized the importance of collaboration. “I think it is important to realize that even though this is NASA’s IPG workshop, the IPG is part of an entire grid effort throughout the United States, and throughout Europe, and the Asian Pacific Rim – I want to make it really evident that we are doing this all together. Many levels of many different organizations are working toward the same goal of merging efforts into a single grid.”

Although presenters came from various institutions and reported on different topics, there was a central theme and direction to all presentations. “It’s pretty clear that the area of web services is going to be the hot topic this year in grids. There was quite a bit of discussion in all of the talks about what web services are,” said Tony Lisotta, NAS Division IPG task leader.

Among the presenters were: Geoffrey Fox, Indiana University, who discussed the need to define interfaces for web applications to facilitate computer-to-computer interactions in “Peer to Peer Networks and Web Services for a Community Grid”; George Myers, NASA Ames Research Center, outlined recent developments in LaunchPad, a web portal created to enable grid users to submit jobs to IPG resources via any web browser; and Robert Griffin, NASA Glenn Research Center, who discussed interactive online job submission for monitoring and modeling aircraft through the Aviation Safety Program.

One outcome of the workshop: a new collaboration was formed at the workshop between NASA’s IPG and the National Virtual Observatory (NVO) project. The NVO project aims to make images of the night sky captured with powerful telescopesavailable to individual astronomers. NVO would like to provide high-powered computational resources necessary to process the data. It is this sort of collaboration that enables the IPG to grow and expand on its capabilities.

NAS Division senior scientist Tom Hinke summed up the event: “I thought the workshop was extremely successful. We had good participation from the general grid community this year, and I hope to see more applications people at future workshops.”

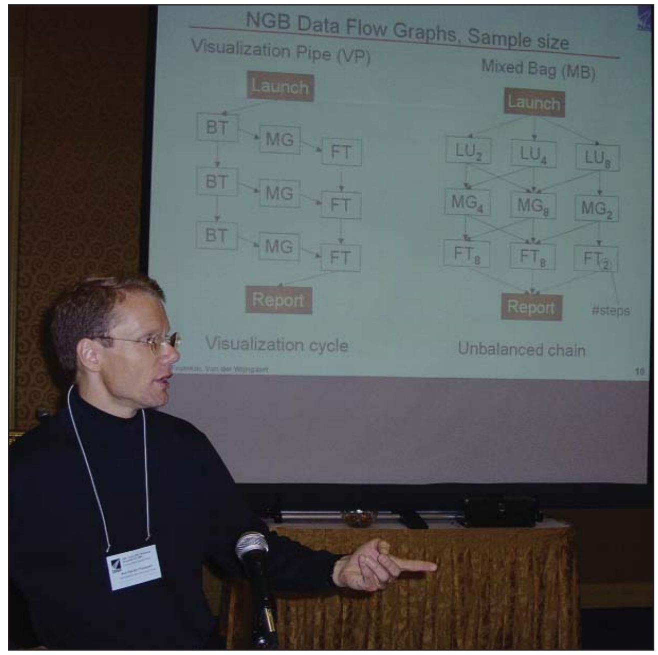


*Rob F. Van Der Wijngaart, NAS Division computer scientist, presented his research on Grid benchmarking at the third annual IPG workshop held in December. (N. Veronico)*

## CFD Cycle Time Flies with Pegasus5 Software

Computational fluid dynamics (CFD) cycle time just got shorter with the newly improved Pegasus5 code. Co-developer Stuart Rogers, an aerospace engineer in the NAS Division, recently applied this code to a difficult and complex configuration – a complete Boeing 777-200 high-lift aircraft.

Starting with existing volume grids, Rogers was able to construct the entire overset grid-system in only three days – a calculation that took 32 days to accomplish in 1998 using the previous version of the software, Pegasus4.

“This speed-up in CFD cycle time and reduction in user expertise requirements will significantly reduce the cost of applying viscous CFD methods to complex design and analysis problems,” Rogers said.

The Pegasus5 software is used to perform the preprocessing task of linking together a large number of randomly overset grids. All that is required to use Pegasus5 is the OVERFLOW input file and the volume grids. During the past year, significant enhancements were made to the software, including algorithm improvements, bug fixes, and parallelization. These advances enable Pegasus5 to perform calculations days sooner than Pegasus4.

Rogers plans to continue enhancements on the most recent version of Pegasus, making additional algorithm improvements to increase the code’s automation and robustness.

# Launching Jobs Into Grid Space

**Information Power Grid users can now send jobs to any IPG resource through the recently developed grid portal, LaunchPad.**

Running jobs using the distributed resources of NASA's Information Power Grid (IPG) can be accomplished with an easy-to-use job-launching tool. Working in conjunction with developers at Lawrence Berkeley National Laboratory (LBNL), researchers in the NASA Advanced Supercomputing (NAS) Division's portal development group designed the IPG LaunchPad for the grid. LaunchPad enables users to create and submit jobs, track job progress, and select resources – all from a web browser. This web-based interface, or portal, also enables researchers to transfer files remotely and, with limited capability, customize their view of the tool to meet individual needs.

"The vision for LaunchPad is that it will enable NASA scientists and engineers to access remote IPG resources, including computers, data archives, instruments, and software resources – all with a single log-in from any desktop computer," says Arsi Vaziri, IPG deputy project manager. IPG LaunchPad was designed using the Grid Portal Development Kit developed by Jason Novotny, formerly with the National Computational Science Alliance, now at LBL.

The job launching tool consists of easy-to-use windows, which prompt researchers to input information about their jobs. LaunchPad was designed with users in mind – it requires a minimal amount of password entries, reducing a researcher's workload while also maintaining security.

## Just Right for Users

"IPG LaunchPad is an evolutionary tool and we plan to make enhancements as needed. For example, examining the environment and automatically running a job on the resource where it will run soonest, or providing a shell-like environment that allows the user to run simple commands such as familiar shell commands without having to wait in a batch queue," explains IPG team member George Myers. LaunchPad is still in the formative stage – the team is working to make it more intuitive. The LaunchPad tool is currently capable of running batch jobs on the grid, securely transferring files, displaying current job status information, and, to a minor extent, creating a customizable environment. Developing a tool that can be fully modified to fit user needs will require a user profile database to store each person's preferences. LaunchPad also enables users to view old output and previously run jobs.

*Continued on page 5*

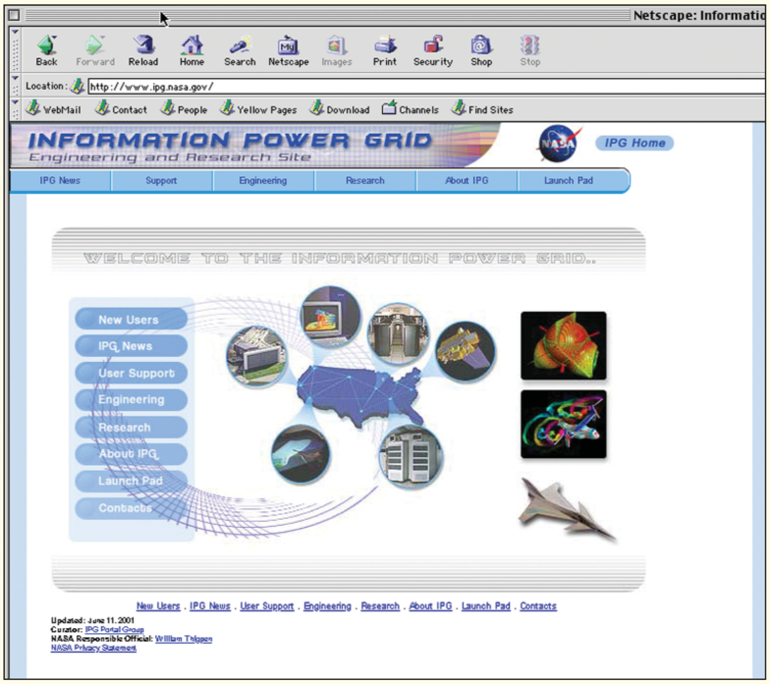


*The Information Power Grid's (IPG) homepage provides users with recent news about the grid effort, information for users new to the grid environment, statistics about the resources available on the grid, frequently asked questions, IPG user support, and much more. A link to the LaunchPad grid portal is located in the upper right corner of the IPG website. To access the job launching tool, users must have a valid certificate.*

# The Evolution of High-Performance Computers in the NASA Advanced Supercomputing Division

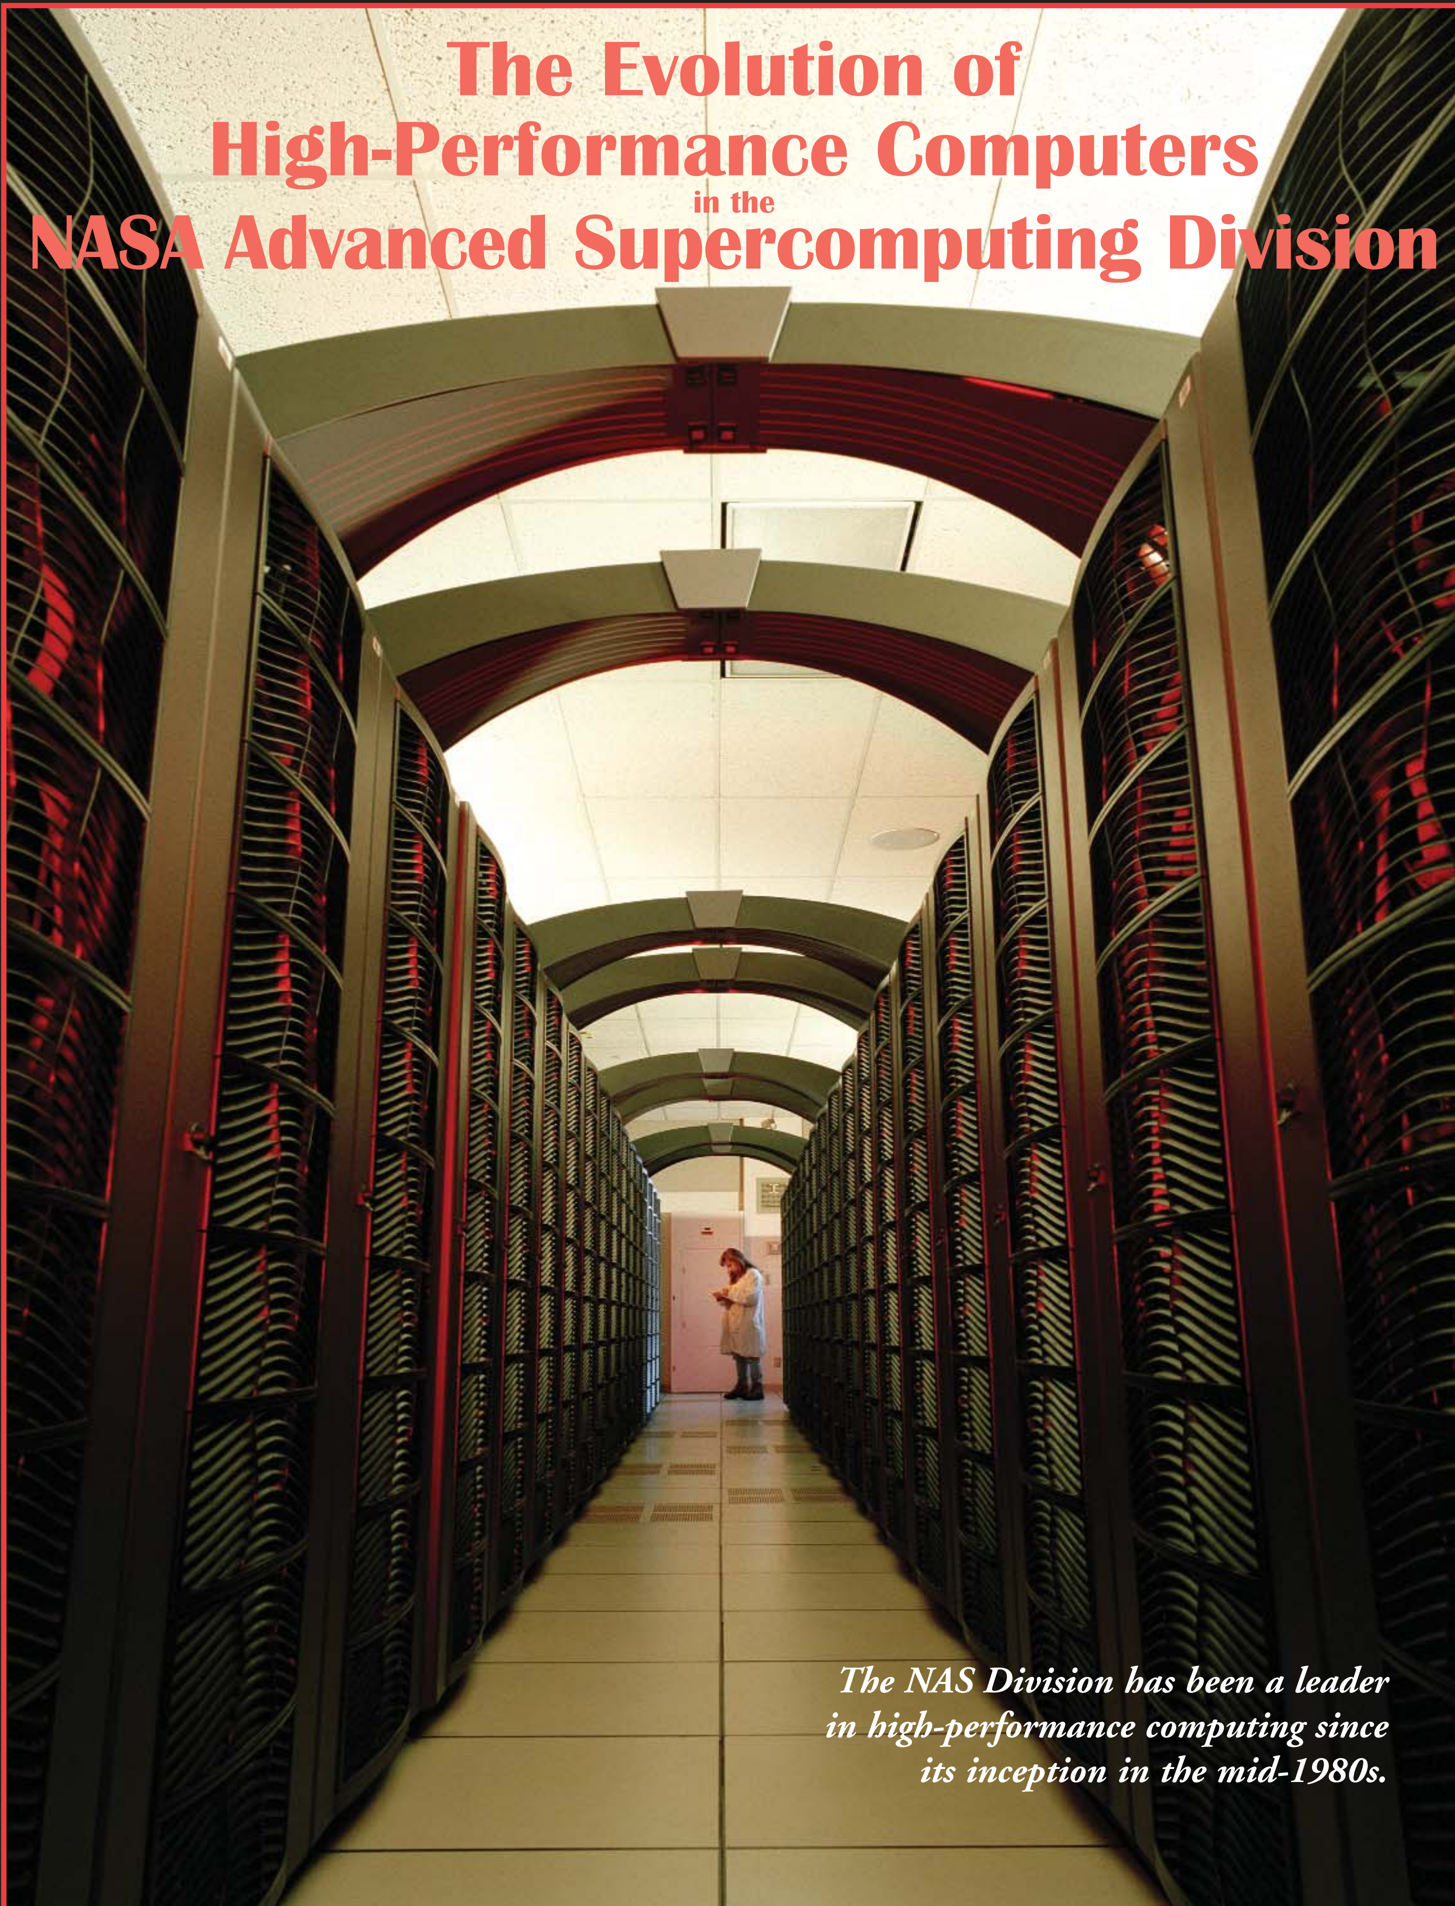

*The NAS Division has been a leader in high-performance computing since its inception in the mid-1980s.*

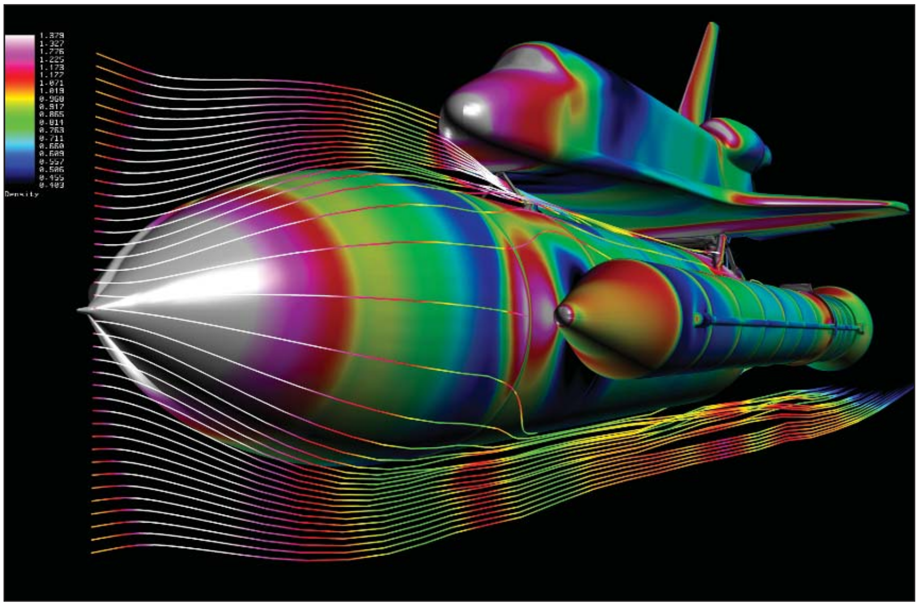

*NASA scientists believed that computers could more accurately represent fluid flows over an air- or spacecraft without the inherent flow interruptions of the wind tunnel's walls and the model support (below). Also, computers would enable rapid design changes without the cost of building a new model. At left, streamlines around a high-resolution (16 million point), 113 zone model of the Space Shuttle Launch Vehicle (SSLV) depict air flow around and between the craft's bodies. A solution of this resolution required 110 hours on the Cray C90. (Left: dataset courtesy Reynaldo J. Gomez/NASA Johnson Space Center; Visualization by Timothy A. Sandstrom/ NASA Ames using GEL software. Below: NASA/Ames)*

The vision for what is today the NASA Advanced Supercomputing (NAS) Division began to materialize in the late 1970s. Scientists at NASA Ames Research Center were pioneering computational fluid dynamics (CFD) under the leadership of Hans Mark, then Ames center director. Mark had witnessed the dramatic advances in nuclear reactor design enabled by computers, and those beginning to be applied to climate modeling, while working at Lawrence Radiation Laboratory. He strongly believed that the promise of computers could be used to develop what he called a "numerical wind tunnel" to solve such problems as airflow over an air- or spacecraft in a variety of flight attitudes, using Navier-Stokes equations in three dimensions.

From Mark's numerical wind tunnel concept, NAS scientists have worked with industry to develop supercomputers capable of solving unsteady 3-D Navier-Stokes equations as well as complex computations in aerospace, nanotechnology, datamining problems, computational chemistry interactions, and challenges in climate simulations and turbulence modeling.

In 1978 a cadre of Ames researchers, F. Ron Bailey, K.G. Stevens Jr., Wayne Hathaway, and Ray Lim, were tasked with formulating the requirements for the numerical wind tunnel concept, selecting computer architectures to solve 3-D Navier-Stokes equations, and selling the idea to Congress, NASA, and the aerospace industry. "CFD was looked at as a way of getting design information that you could not get, or could not get efficiently, in a wind tunnel," recalls Bailey, NAS's division chief from its inception through July 1990. "Of particular interest was transsonic flow. In the military sense, tactics had changed from the 1950s idea of flying supersonically to attack one's enemy with a missile," Bailey explains. Vietnam showed that pilots had to slow down to transsonic speeds to maneuver and attack. Thus, the military needed aircraft that were more maneuverable and flew better at transsonic speeds where the actual combat occurred. At the same time, jet transport aircraft were bumping up against the transsonic drag rise (at Mach .85 aircraft develop local

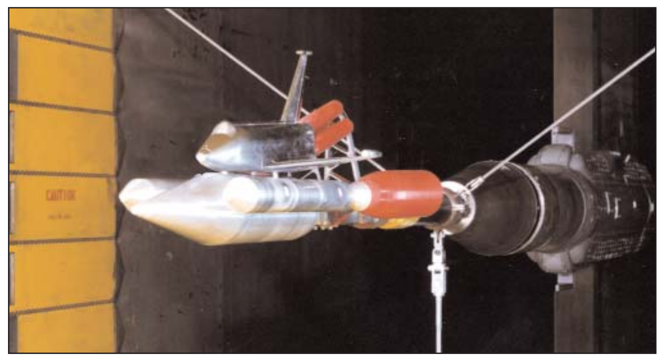

regions of supersonic flow that causes increased drag), which happens to be the optimum range for a given fuel load and weight.

"It is very difficult to get transsonic design data, and it is difficult to obtain in the wind tunnel because you have ventilation – and the tunnel's walls as well as the model support generate interference. In addition, Navier-Stokes equations are mathematically challenging because they are very strongly non-linear problems. Mathematicians just simply could not solve the problem, so it had to be solved with computers," Bailey says. The vision was to use computers capable of calculating CFD datasets to enable researchers to make a variety of changes to an aerospace vehicle's design quickly and cheaply. Once a favorable design was selected, a model could be built and used to validate the design's performance in a wind tunnel.

Researchers developing the NAS Program envisioned a single facility housing supercomputers, the scientists using the machine, and support personnel. This facility would serve as the hub for the high-performance computing community of NASA, the Department of Defense, universities, and the U.S. aerospace industry.

"From 1976 to 1979, a series of studies were conducted producing baseline configurations, a functional design, and rough estimates of cost and schedule for the NAS program," says Bailey. “In September 1980, two parallel design-definition contracts were awarded, and almost two years later, in April 1982, the contractors’ proposals for detailed design, development, and construction were submitted for evaluation. It soon became clear that the earlier approaches needed to be changed. First, computational aerodynamics had advanced rapidly, and it was deemed necessary to not only establish, but to sustain a state-of-the-art computational aerodynamics facility. Also, supercomputers had advanced to the point that it was no longer considered necessary to directly subsidize their future development.

“As a result,” Bailey continues, “a new NAS program plan was developed in February 1983. The NAS Program was defined by three principal goals: to provide a national computational capability as a necessary element in ensuring continuing leadership in CFD and related disciplines; to act as a pathfinder in advanced, large-scale computer system capability through the incorporation of state-of-the-art improvements in computer hardware and software technologies; and to provide a powerful research tool for NASA’s Office of Aeronautics and Space Technology.”

## NAS Takes Shape

The core of the NAS Division’s researchers came from the CFD branch at Ames. Operations began with the acquisition of a Cray XMP-12 (12 standing for one CPU – central processing unit, and two Megawords of internal memory), delivered in August 1984. This Cray was used for initial system integration and to test the Unix operating system on a supercomputer. The X-MP at NAS was the first Unix-based Cray supercomputer delivered by the company, and was the first to be interfaced with Unix workstations. It enabled researchers to compute low-resolution, steady-state 3-D problems such as air flow over an aircraft’s wings in various flight attitudes.

In April 1985, NAS partnered with Cray Research Inc. and the National Magnetic Fusion Energy Computer Center (NMFECC), Livermore, Calif., to acquire a single CPU Cray-2. The machine was located at NMFECC, and each partner had eight hours of daily access.

The division acquired its first Cray 2, serial number 2002, on September 30, 1985. The machine was housed at the Ames Computer Processing Facility. “This was the first Cray 2 delivered with a 256-megaword memory, which was larger than the memory installed in all previous Cray machines delivered – combined,” Bailey says.

To honor the pioneers of science and aerospace, division management began the tradition of naming its computers – large machines after mathematicians and physicists, and smaller machines after prominent figures in aviation. Thus, Cray 2, serial number 2002, was named *Navier*, in honor of Claude Louis Marie Henri Navier (1785– 1836), a French professor of applied mechanics who co-developed the Navier-Stokes equations of fluid dynamics for incompressible fluids and viscous fluids.

Users across the United States were linked to the NAS Division’s computers through a high-speed, long-haul communications subsystem known as NASnet. This network became operational in July 1985, using the data transfer protocol TCP/IP (transmission control protocol/Internet protocol), which later became the standard that enabled the development of the Internet. NASA’s Dryden (California), Langley (Virginia), and Glenn (Ohio) Research Centers, industry, academia, and numerous government agencies connected to NASnet through the Department of Defense’s MILnet (military network) and ARPAnet (Advanced Research Projects

# NAS Computer Names

*The NAS Division honors its scientific heritage by naming its high-performance computers after pioneering people in aviation, mathematics, space, and computing. The following figures, both contemporary and historic, and their namesake machines are listed below.*

*Amelia* (VAX 11/780)
Named for pioneering aviatrix Amelia Earhart (1897–1937) who was the first woman and the second person to fly solo across the Atlantic on the fifth anniversary of Charles Lindbergh’s historic 1927 crossing.

*Babbage* (IBM SP-2)
Charles Babbage (1791–1871) is considered the grandfather of the digital computer. He believed that a computing machine must be composed of an input device (he used a card reader), a memory device (he called “The Store”), a central processing unit (he called “The Mill”) and an output device (he used a printer).

*BJ* (DAO – SGI Origin 2000)
Vilhelm Frimann Koren Bjerknes (1862–1951) was a Norwegian physicist and pioneer in modern meteorology. He worked on applying hydrodynamic and thermodynamic theories toward predicting future weather conditions. His work in meteorology and on electric waves was important in the early development of wireless telegraphy. Bjerknes developed a theory of cyclones, known as the polar front theory, with his son Jakob Aall Bonnevie Bjerknes (1897-1975).

*Boltzmann* (Thinking Machine CM5)
Ludwig Eduard Boltzmann (1844–1906) Austrian physicist known for his work in thermodynamics and on the kinetic theory of gases.

*Bright* (Cray SV-1)
Loren G. Bright (1924–) joined the National Advisory Committee for Aeronautics, the precursor to NASA, in 1947, and served as Director of Research Support at Ames from 1969 until his retirement in 1979. During his nine-year tenure as a director, Bright managed computer systems development and operations; manufacturing; physical plant maintenance; and the construction of new research

*Continued on page 4A*

*Continued from page 3A*
facilities at the center. As an early advocate for supercomputing's potential benefit to NASA, Bright worked through the agency's Institute for Advanced Computation to develop and secure funding for both near- and long-term research programs at Ames. Bright was awarded the NASA Outstanding Leadership Medal in 1978.

*Chapman* (SGI Origin 3800)
Under Dr. Dean Chapman's (1922–1995) guidance as director of Ames' Thermo and Gas Dynamics Division (1969), and later Director of Astronautics (1974), Ames developed thermal protection systems for the Space Shuttle, Galileo Jupiter probe, and many other vehicles. His creation of the Computational Fluid Dynamics Branch at Ames in 1970, using the best available high-speed computers at that time, assured the center's preeminence in that new and exciting field of research. The first "massively parallel" computer of its time, the Illiac IV, was installed at Ames in 1970 and, under Chapman's leadership, became operational in 1973.

*Chuck* (Convex 3820)
Brig. Gen. Charles E. "Chuck" Yeager (USAF Retired, 1923–) logged more than 10,000 hours in 180 different military aircraft, including foreign and experimental rocket aircraft. In 1947, Yeager exceeded Mach 1 at Edwards Air Force Base, California, piloting the X-1.

*da Vinci* (SGI Power Challenge XL)
This workstation cluster was named for sculptor, painter, architect, engineer, and scientist Leonardo da Vinci (1452–1519).

*Dixon0* (DAO – SGI Origin 2000)
Dixon Butler is the former NASA EOS Project scientist and a major advocate for supercomputing and data servicing in Earth Sciences. He became chief scientist of the GLOBE (Global Learning and Observations to Benefit the Environment) Program.

*Eagle* (Cray C90)
Named for the first manned space capsule to land on the Moon. "The Eagle has landed." These words ushered in a new era of human exploration on July 20, 1969, as the first
*Continued on page 5A*

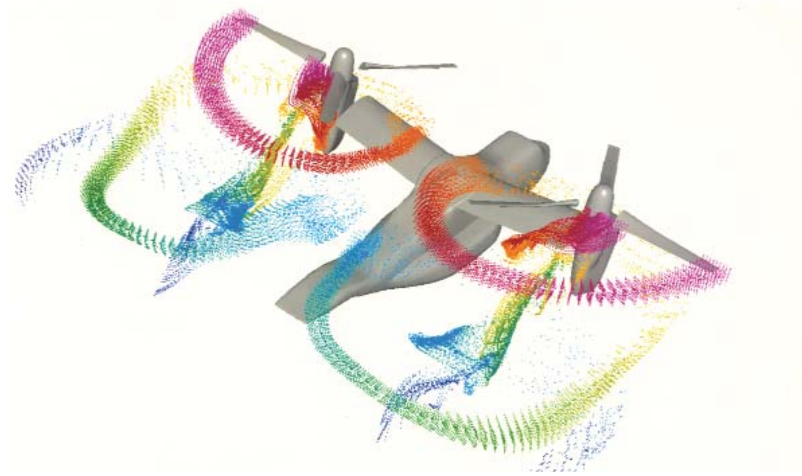

Agency network), NSFnet (National Science Foundation network), BARRnet (Bay Area Regional Research network connecting University of California campuses at Berkeley, Davis, San Francisco, and Santa Cruz, plus Stanford University), as well as the San Diego Supercomputer Center. NASnet brought the power of supercomputers to researchers in the fields of aerodynamics, astrophysics, life sciences, turbulence modeling, as well as computational fluid dynamics and computational chemistry across the United States.

During its first eight months of operation (July 1986 through February 1987), NAS's Cray 2 supported 123 projects: 30 percent were devoted to research for the National AeroSpace Plane (a Reagan-era conceptual airliner/spacecraft powered by hydrogen burning supersonic-combustion ramjets), 42 percent to aerodynamics, and 11 percent to turbulence modeling. The remaining 17 percent was dedicated to astrophysics, chemistry, atmospheric, and life sciences. When accessing the Cray 2's entire 256-megaword memory to solve a 3-D Reynolds-averaged Navier Stokes equation using in-core memory, the machine could calculate 10.24 million floating points in 24 hours. This was an impressive number for its day – in comparison, the NAS Division's high-performance computers can now calculate more than 200 million points per day.

## NAS Becomes a Production Facility

The vision of the numerical wind tunnel became a reality on March 1, 1987. A 90,000 square-foot facility was built at Ames Research Center to house the division's supercomputers, researchers, and support personnel. This facility was constructed on the north side of Ames, and began operations on March 1, 1987. Eight days later, then NASA Administrator James C. Fletcher dedicated the new NAS Facility, which housed the Cray 2 *Navier*, and the newly delivered Thinking Machines CM2, *Pierre*. The CM-2 was an experimental, single instruction, multiple data machine employed by the NAS advanced development laboratory to study the interaction between high-speed parallel processors.

The following January, a second Cray 2, *Stokes*, was installed at the NAS Facility – essentially doubling production computing power. In addition to providing researchers with more capability, both Cray 2s were used to develop and test a variety of new supercomputer software. NAS computer scientists also established the NAS Parallel Benchmarks (NPB), which became the industry's standard for measuring the performance of parallel computing architectures.

On August 30, 1988, the very first Cray Y-MP/832 (8 CPUs and 32 Megawords of internal memory), *Reynolds*, was delivered to NAS. After initial testing and a pre-production phase, *Reynolds* was added to the produc-

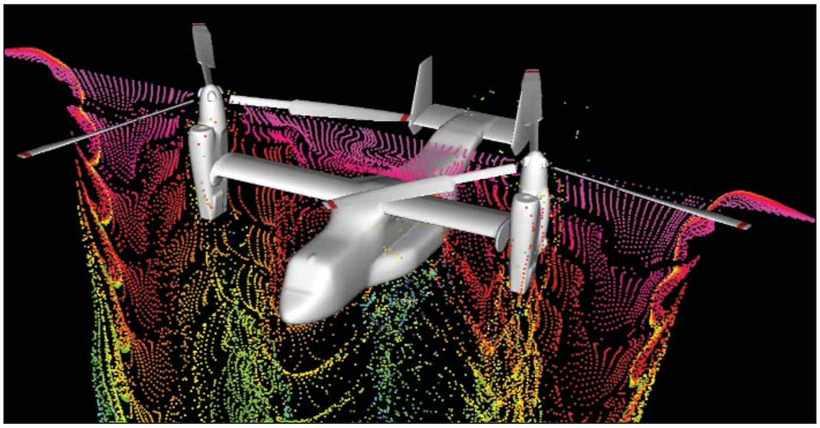

*The evolution of high-performance computers at NAS has increased the fidelity of computational fluid dynamics modeling as evidenced by the 1993 particle tracing of two blade revolutions of a tilt-rotor aircraft (opposite page) to the particle traces around a large-scale unsteady moving body. (Opposite page: Robert Meakin and David Kao. Above: M. Jahed Djomehri)*

tion pool on February 15, 1989. *Reynolds*, in comparison, could solve an unsteady 3-D Navier-Stokes equation 60 percent faster than the three-year-old Cray 2 *Navier*.

Two Cray C90s (*von Neumann* and *Eagle*) arrived at NAS in the spring and early summer of 1993. In February 1996, a Cray J90, *Newton*, was added to the NAS production pool. These machines provided the majority of the production computing cycles for the remainder of the decade. *Von Neumann*, installed in March 1993, was operated by the NAS Division until March 2002. Operating and maintenance costs, as well as a shortage of spare parts, drove the decision to retire the machine, which was replaced by a Cray SV1ex.

## Experimental Machines Play an Important Role

While the Cray machines worked day-in and day-out to provide computing cycles to researchers, the advanced development team was experimenting with an Intel Touchstone iPSC/860, the Thinking Machines Corporation's massively parallel CM5, and the Intel Paragon, which offered improved interconnects when compared to the iPSC860. An IBM SP-2, *Babbage*, joined the experimental machine pool during the summer of 1994. It was similar to the Thinking Machines CM5 architecture, where each processor operated on its own with a separate set of instructions. But administration and upgrading of the SP-2 was always a challenge, and the machine's complexity, coupled with the advances of other supercomputers, quickly made this machine obsolete.

Programming the majority of these experimental systems was extremely complex, making them unsuitable for the production computing environment. In addition, their computational speeds were constrained, and their compilers were not compatible with production-quality codes.

The evolution of supercomputers from 1975 to 1995 has been compared to the early days of the automobile industry, which saw cars with three wheels; combustion engines; steam engines; driver on the right, in the middle, on the left; hard tops, soft tops, no tops. It took one company to set the standard. Once that was accomplished, the industry moved forward and those companies with inferior products left the marketplace. Similarly, the high-performance computing industry saw a variety of architectures and more than a dozen companies competing for a small client base. Some vendors

*Continued on page 8A*

*Continued from page 4A*

manned flight to the moon touched down in the Sea of Tranquility. "That's one small step for man, one giant leap for mankind." – Neil A. Armstrong

*Evelyn* (SGI Origin 2000)
Evelyn Boyd Granville (1924–) was one of the first African-American women to earn a doctorate in Mathematics. In the early 1960s, she developed computer programs that were used for trajectory analysis in the Mercury Project (the first U.S. manned space mission), as well as the Apollo Project (to explore the Moon).

*Fred* (VAX 11/780)
Frederick J. "Fred" Noonan (1893–1937) was one of the most outstanding navigators of his time, and in the mid-1930s certainly one of the most experienced aerial navigators. He charted the courses for Pan Am exploratory flights across the Pacific Ocean to Hawaii, Midway Island, Wake, Guam, the Philippines, and Hong Kong. Noonan served as navigator on Amelia Earhart's 1937 around-the-world flight. Both disappeared between Lae, New Guinea, and Howland Island.

*Fermi* (SGI Origin 2000)
Enrico Fermi (1901–1954) was the first physicist to split the atom, and his research was instrumental in the development of nuclear power generation. Fermi was awarded the Nobel Prize in 1938 for his nuclear research. (Fermi is a set of CPUs within *Turing*.)

*Gamma* (Intel iPSC/860)
The third letter of the Greek alphabet, as well as a unit of magnetic field intensity (10-5).

*Gatun* (DAO – SGI Origin 2000)
Gatun was named after a lock in the Panama Canal. Gatun stores all the input observations (about 400,000 per day) and the DAO operational output data products (about 3.5 gigabytes per day with GEOS DAS v3.0).

*Gottfried* (VAX 6320)
Gottfried Wilhelm Leibniz (1646–1716) invented differential and integral calculus simultaneously with Newton.

*Continued on page 8A*

# NASA Advanced Supercomputing (NAS) Division comp

1987 • *Scott/Chuck*
Convex 3820
Peak MFLOPS/processor: 240

1984 • Cray X-MP
Peak MFLOPS/processor: 210.53

1987 • *Pierre*
Thinking Machines CM-2
Peak MFLOPS/processor: 14

1990 • *Gamma*
Intel iPSC/860
Peak MFLOPS/processor: 60

1986 • *Navier* (*Stokes* pictured)
Cray 2
Peak MFLOPS/processor: 487.8

1988 • *Reynolds*
Cray Y-MP
Peak MFLOPS/processor: 317.46

1
Thin
Peak M

Peak

1

Peak

1983 1984 1985 1986 1987 1988 1989 1990 1991 1992

serial

vector

parallel

# uter Timeline

www.nas.nasa.gov
www.nas.nasa.gov/gridpoints

993 • *Boltzmann*
king Machines CM-5
FLOPS/processor: 128

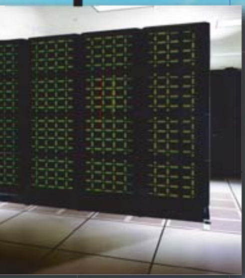

1993 • *Sigma*
Intel Paragon
k MFLOPS/processor: 75

993 • *Von Neumann*
Cray C90
MFLOPS/processor: 960

1994 • *Babbage*
IBM SP-2
Peak MFLOPS/processor: 266

1996 • *Newton*
Cray J90
Peak MFLOPS/processor: 200

1997 • *Turing*
SGI Orgin 2000/195 Mhz
Peak MFLOPS/processor: 390

1999 • *Lomax*
SGI Orgin 2800/400 Mhz
Peak MFLOPS/processor: 800

2001 • *Bright*
Cray SV1ex
Peak MFLOPS/processor: 2,000

2001 • *Chapman*
SGI Orgin 3800/400 Mhz
Peak MFLOPS/processor: 800

1993 1994 1995 1996 1997 1998 1999 2000 2001 2002

Distributed shared memory/Distributed computing

*Continued from page 5A*
*Helios1* and *Helios2* (DAO – SGI Origin 2000)
The Sun god, son of Hyperion: later identified with Apollo.

*Hopper* (SGI Origin 2000)
Rear Admiral Grace Murray Hopper (1906–1992) was a mathematician and pioneer in data processing. The admiral made several vital contributions to the development of modern computing systems in the early 1960s, including helping to invent the Cobol programming language as well as the first practical compiler for modern computers.

*Jimpf0* and *Jimpf1* (DAO – SGI Origin 2000)
Jim Pfaendtner, a noted computer scientist at Goddard Space Flight Center, is currently chief of the Production Management Branch of the National Meteorological Center and is responsible for the computer activities associated with daily weather forecasting for the nation.

*Kalnay* (DAO – SGI Origin 2000)
Eugenia Kalnay is the chair of the Meteorological Department at the University of Maryland and was director of the Environmental Modeling Center for the National Centers for Environmental Prediction (NCEP), National Weather Service.

*Lagrange* (Intel 386)
Joseph Louis Lagrange (1736–1813) was an 18th century mathematician most famous for his theory that expresses a dynamic system through the function of coordinates, velocities, and times.

*Lomax* (SGI Origin 2800)
Harvard Lomax (1922–1999) was a scientist and manager at NASA Ames Research Center for more than 50 years. He was instrumental in leading the transition from theoretical aerodynamics to computational aerodynamics, and personally performed pioneering research in theoretical aerodynamics, numerical algorithm development, parallel computing, computational turbulence, and turbulence modeling.

*Lou* (DAO – SGI Origin 2000)
Louis Lopez (1938–1999) was a scientist in the NAS Division. He began his professional career with IBM at their Scientific Center in Houston, where he studied computer simula-
*Continued on page 9A*

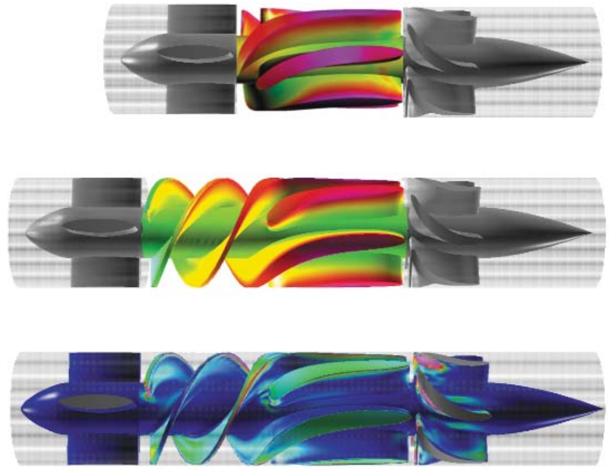

*NASA computational fluid dynamics researchers applied Space Shuttle rocket fuel pump technology to the DeBakey Heart Assist Device. Compare the original ventricular assist device (top) to the unit after modifications by NAS researchers (bottom). NAS Division researchers added an inducer to the DeBakey device, which eliminated the dangerous backflow of blood by increasing pressure and making the flow more continuous. The device is subject to the highest pressure around the blade tips, shown in magenta. NASA recently named the DeBakey Heart Assist Device as its Invention of the Year for 2001. (NASA/Cetin Kiris)*

*Continued from page 5A*
consolidated with other, larger companies, while many simply disappeared. The majority of these computers were better at contrasting datasets (cryptographic computing for the intelligence community, for example) while others, notably the Crays, were superior at processing large blocks of arithmetic (as used in the NAS Division's computational physics). However, these experimental machines did provide a powerful platform for developing new machine architectures to advance the state-of-the-art in high-performance computing and to push for architectures that could compute at Teraflop speed.

NAS researchers used the pool of experimental machines to develop the Portable Batch System (PBS). Developed collaboratively with computer scientists at Lawrence Livermore National Laboratory, the software was introduced on March 1, 1994, to synchronize job scheduling and to run batch jobs from workstations. PBS was eventually transferred to the commercial firm MRJ/Veridian, which now distributes the software.

## 1,024 and Beyond

In May 1996, then NASA Administrator Daniel S. Goldin tasked Ames Research Center with becoming the agency's Center of Excellence for Information Technology. Ames took the lead in fostering research and development relationships with the most promising high-performance computer vendors. Out of these pathfinding relationships (with vendors including Cray, Intel, IBM, SGI, and Thinking Machines), SGI's Origin 2000 was the only system able to sustain a level of performance on 3-D Navier-Stokes CFD problems exceeding the sustained performance of the 16-processor Cray C-90. SGI's new offering featured a non-uniform memory access (NUMA, later ccNUMA – for cache-coherent NUMA) multiprocessor that would eventually scale to more than 1,000 processors with minimal performance (bandwidth or latency) penalties.

NAS researchers were excited about the scientific possibilities the Origin computers offered, and within four months a 64-CPU Origin 2000 system, *Hopper,* was delivered and installed. *Hopper* was used to develop technologies

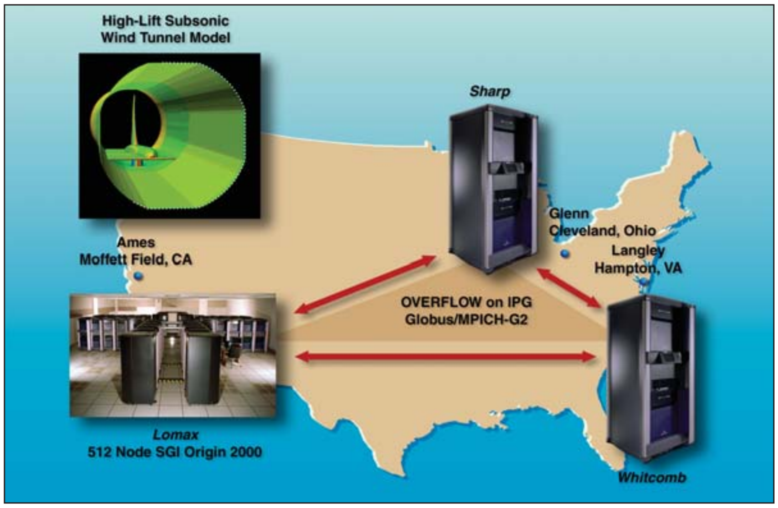


*NASA's Information Power Grid (IPG) enables researchers to use computers in both a parallel and distributed fashion. The IPG can manage underutilized resources, decrease computational turnaround time, support a collaborative problem-solving environment, and provide more cost-effective computing solutions. In this example, researchers solved a complex 14-million-point aerospace geometry calculation (high-lift subsonic wind tunnel data) using three remotely located computers at Ames, Glenn, and Langley Research Centers.*

for the future 256-processor Origin computer. When the CFD code ARC3D was ported to this machine, it demonstrated a sustained speed of 20 gigaflops – a five time improvement over the 16-processor Cray C90 *von Neumann.* (Many of the Origin systems that reside in the NAS Facility were jointly acquired by NASA's Data Assimilation Office [DAO], the Consolidated Supercomputing Management Office [CoSMO] and Ames Research Center's Information Technology Program.) In addition to the NUMA/ccNUMA architecture, these computers offered increased computing power at a lower acquisition cost. The Origins set the stage for the NAS Division's foray into scientific areas outside of aeronautics, while at the same time capturing the essence of Administrator Goldin's vision.

As government had driven the development of computers in previous decades, NAS computer scientists collaborated with SGI to build a 256-processor machine incorporating single system image (SSI) technology, which uses a single operating system to control all the processors attached to a globally shared memory as well as the input/output system. The success of the 256-processor SSI Origin 2000, *Steger*, enabled NASA to fund an advanced 512-processor Origin. SGI and NAS worked together to minimize the technical risk in developing a 512-processor Origin 2800, known as *Lomax*. After the success of the high-processor Origins at NAS, SGI was able to spin off the machine designs and found a commercial market for both the 256 and later 512 systems.

The performance improvement of NAS's computational codes on the 512-processor Origin 2800 attracted the attention of NASA as well as the scientific and computer development communities. NAS researchers and SGI computer scientists believed that pairing two newer technology 512-processor machines could increase performance by a factor of three.

On August 23, 2000, NAS and SGI jointly announced plans to build the first 1,024 processor supercomputer using the SSI architecture. Named *Chapman*, the computer began operations in August 2001. "The previous 512-processor Origin 2800 showed that CFD and climate modeling codes scaled linearly, and that a 1,024-processor Origin 3800 would likely do the

*Continued from page 8A*

tions of blood flow. Later, he began working on supercomputer performance analysis tools, first at IBM's Palo Alto Scientific Center and then at Ames. His efforts on improving parallel program performance resulted in the NAS Trace Visualizer (NTV), a tool that helps users understand what happens during a program's execution.

*Navier* (Cray 2)
Claude Louis Marie Henri Navier (1785–1836) was a French professor of applied mechanics at the Ecole des Ponts et Chaussées beginning in 1830. Navier is remembered for the Navier-Stokes equations of fluid dynamics for incompressible fluids and viscous fluids.

*Newton* (Cray J90)
Sir Isaac Newton (1643–1727), an English mathematician and physicist, was considered one of the greatest scientists in history. Newton was one of the inventors of the branch of mathematics known as Calculus. Newton solved the mysteries of light and optics, formulated the three laws of motion, and derived from them the law of universal gravitation.

*Orville* (VAX 11/780)
Orville Wright (1871–1948) and his brother Wilbur flew the first powered aircraft in December 1903. The plane was built in Dayton, Ohio, and shipped to Kitty Hawk on the North Carolina coast for its first flight.

*Pierre* (Thinking Machines CM-2)
Pierre-Simon LaPlace (1749 – 1827) was a mathematician and physicist who proved the stability of celestial bodies through differential equations.

*Piglet* (SGI Origin 2000)
Named for the fictional character in A.A. Milne's *Winnie The Pooh.*

*Poncho* (1901–1975)
Florence L. "Poncho" Barnes was an aviatrix and proprietor of the Rancho Oro Verde Fly-Inn Dude Ranch on the outskirts of what is now Edwards Air Force Base, California. Poncho was friend and confidant to America's pioneering aviators including Charles "Chuck" Yeager and James H. "Jimmy" Doolittle.

*Continued on page 10A*

*Continued from page 9A*

*Prandtl* (Amdahl 5880 – NAStore)
Ludwig Prandtl (1875–1953), mathematician and professor of applied mechanics at the University of Gottingen (Germany) where he founded the school of aerodynamics. Prandtl discovered the boundary layer of friction and drag over a wing in 1904.

*Reynolds* (Cray Y-MP)
Osborne Reynolds (1842–1912) was the first Professor of Engineering at the University of Manchester, England, who produced many of the modern theories of turbulent flow, hydraulic modeling, hydrodynamic lubrication, friction, and heat transfer. Known for the "Reynolds Number" used in fluid dynamics calculations.

*Scott* (Convex 3820)
Scott Crossfield was the first person to reach Mach 2.005 at Edwards Air Force Base, Calif., in 1953 piloting a Douglas D-558-II Skyrocket. Crossfield flew every aircraft in the nation's inventory. He flew the X-15, an aircraft he helped design, and was the first to fly the X-15 in a glide flight and a powered flight.

*Sigma* (Intel Paragon)
Sigma is the 18th letter of the Greek alphabet.

*Siméon* (Sun 4/490)
Siméon Denis Poisson (1781–1840) developed the theories of electrostatics and magnetism.

*Simak* (Sun Enterprise 10000)
Clifford D. Simak (1904–1988), was born in Minneapolis, Minn., and attended the University of Wisconsin. Simak worked with various newspapers in the Midwest from the early 1930s to 1961. He is considered one of the elders of science fiction writing. Often described as a "pastoralist" for presenting America's small-town viewpoint. His focus was on human philosophy, the countryside, and people.

*Steger* (SGI Origin 2000)
Joseph Steger (1943–1992) began his career as a scientist at NASA Ames Research Center, and later became a professor at Stanford University and the University of California at Davis. He pioneered computational technology, both

*Continued on page 11A*

same," says NAS Senior Scientist Jim Taft. "We demonstrated that this was true within a matter of weeks. The OVERFLOW code, used in CFD, scaled from 60 gigaflops on the 512 to 165 gigaflops on the 1,024, while the climate modeling code (FVCCM3-MLP) scaled from 900 days a day to more than 2,900."

Simultaneous to the development of the Origin high-performance computers, NASA began building a network of high-performance computers, data storage archives, and scientific instruments that allows remotely located users access to high-performance computing resources. Known as the Information Power Grid (IPG), its development is being directed by researchers at NAS in collaboration with scientists at Glenn and Langley Research Centers, as well as the Jet Propulsion Laboratory in Pasadena. "The IPG is becoming the model by which NASA accesses its high-performance computing resources," says IPG Project Manager William E. Johnston. "We felt we had an opportunity to fundamentally change the way computing is done – and the IPG accomplishes that." The addition of the 1,024-processor machine further increases the power of the nodes on the IPG. "One of NASA's great strengths is to carry research tasks through to working systems," says Bill Feiereisen, former NAS division chief. "This has been the IPG team's focus for the last three years – to put in place an infrastructure to provide researchers and engineers the computing power to advance science. A working grid now exists that includes high-performance systems at three NASA research centers. In addition, we are about to link IPG with the grids that our colleagues are building at the National Science Foundation/Partnerships for Advanced Computational Infrastructure centers."

"The future of high-performance computing will see faster, higher processor count supercomputers serving researchers located hundreds or thousands of miles from where the machines are located," says NAS Division Chief John Ziebarth. "Our future lies in delivering more computing cycles to the scientific community to enable NASA to meet its mission goals."

*—Nicholas A. Veronico*

## NAS Logo Displays Division Mission

NAS

NUMERICAL AERODYNAMIC SIMULATION

NAS

The NAS Division's logo, as well as its name, has evolved over the years to reflect the division's diverse areas of research. The original logo, 1986, above left, presents the division's focus on computational aerodynamics. At that time the division was known as the Numerical Aerodynamics Simulation Systems Division. A decade later, the logo changed to showcase not only the division's aerodynamics research, but the great strides in turbulence modeling and scientific visualization. In the interim, the name was changed to the Numerical Aerospace Systems Division. The current logo, far right, introduced in 1992, was designed by division employee John Hardman. It illustrates the movement of an aerospace vehicle through a computational grid. The grid represents the NAS Division's work in aero-, nano- and bio-technology, computational chemistry, turbulence modeling, computational visualization, and grid computing. In 2001, the division's name was changed to the NASA Advanced Supercomputing Division.

# Center, Directorate and NAS Division Management

The NAS Division was conceived under the following management structure, and is now a component of the Information Systems Directorate. Management members listed under the Ames center director, director of research support, Thermo and Gas Dynamics Division, and Aerophysics Directorate were instrumental in seeing the vision of NAS Division to fruition.

### NASA Ames Center Directors

Hans Mark (1969–1977) • Clarence Syvertson (1977–1984)
William Ballhaus (1984–1989) • Dale Compton (1989–1994)
Ken K. Munechika (1994–1996)
Henry McDonald (1996 to present)

### Director of Research Support

Loren Bright (1969–1979) • Bob Eddy

### Thermo and Gas Dynamics Division

Victor Peterson

### Aerophysics Directorate

Dean Chapman • William Ballhaus
Victor Peterson • F. Ron Bailey

### Information Systems Directorate

| | |
|---|---|
| David Cooper | July 1994 – April 1997 |
| Henry Lum (Deputy) | Sept. 1994 – April 1997 |
| Steven Zornetzer | April 1997 – present |
| Paul Kutler (Deputy) | April 1997 – Jan. 2001 |

### NAS Division Chiefs

| | |
|---|---|
| F. Ron Bailey | Feb. 1983 – July 1990 |
| Ron Deiss (acting) | July 1990 – Feb. 1991 |
| David Cooper (acting) | March 1991 – Aug. 1991 |
| David Cooper | Aug. 1991 – June 1994 |
| Walter Brooks (div. mgr.) | May 1995 – Nov. 1996 |
| Marisa Chancellor (acting) | Dec. 1996 – Aug. 1997 |
| Paul Kutler (acting) | Aug. 1997 – June 1998 |
| William J. Feiereisen | July 1998 – March 2002 |
| John Ziebarth (acting) | March 2002 – present |

### NAS Deputy Division Chiefs

| | |
|---|---|
| Marcie Smith | 1983 – 1985 |
| Ron Deiss | 1987 – March 1993 |
| Marisa Chancellor | April 1993 – Aug. 1997 |
| Tom Woodrow (acting) | July 1994 – Oct. 1994 |
| Eric A. Hibbard (acting) | Oct. 1994 – Jan. 1997 |
| Tom Lasinski (acting) | Jan. 1997 – Sept. 1999 |
| John P. Ziebarth | Sept. 1999 – March 2002 |

*Continued from page 10A*

theory and application, that revolutionized the use of high-speed digital computers to solve complex problems in aerospace sciences.

*Stokes* (Cray 2)
Sir George Stokes (1819–1903), Irish-born physicist and mathematician, was a pioneer researcher of the motion of incompressible fluids and the friction of fluids in motion.

*Sunrise* (DAO – SGI Origin 2000)
Data Assimilation Office (DAO) cluster used to produce research-quality assimilated global datasets for advancing the understanding of the Earth's system and climate change.

*Turing* (SGI Origin 2000)
Alan Turing (1912–1954) was one of the great pioneers in the computer field. As a mathematician, he applied the concept of the algorithm to digital computers. His research between computers and nature created the field of artificial intelligence. During World War II, code-breaking work under Turing's supervision was very successful thanks to the use of machines that later developed into the first computers.

*von Neumann* (Cray C-90)
John von Neumann (1908–1957) was a brilliant mathematician whose contributions include the bit as a measurement of computer memory; the theory of cellular automata; the serial execution of instructions; and the stored-program concept of computing, a part of virtually every modern computer. "If people do not believe that mathematics is simple, it is only because they do not realize how complicated life is." – von Neumann

*Whitney* (Intel Pentium Pro)
Eli Whitney (1765–1825) was a pioneer mechanical engineer and manufacturer. Whitney is best remembered as the inventor of the cotton gin. He also affected the industrial development of the United States when, in manufacturing muskets for the government, he translated the concept of interchangeable parts, or "commodity" components, into a manufacturing system, giving birth to the American mass-production concept.

*Wilbur* (VAX 11/780)
Wilbur Wright (1867–1912) was co-inventor of the first heavier-than-air powered airplane, which flew at 10:35 a.m., on Dec. 17, 1903.

# Selected NAS Supercomputer Specifications

| Machine Name | Type of Machine | Peak MFLOPS per processor | Number of CPUs | Theoretical Peak Speed Machine GFLOP/second | Main Memory | Installed |
|---|---|---|---|---|---|---|
| ----- | Cray XMP-12 | 210.53 | 1 | 0.21 | 2 Mwords | 08/01/84 |
| *Navier* | Cray 2 | 487.80 | 4 | 1.95 | 256 Mwords | 09/30/85 |
| *Navier** | Cray 2 | 487.80 | 4 | 1.95 | 256 Mwords | 00/00/87 |
| *Chuck/Scott* | Convex 3820 | 240 | 8 | 1.90 | 4 GB | 00/00/87 |
| *Pierre* | Thinking Machines CM2 | 14 | 1,024 | 14.34 | 4 GB | 03/00/87 |
| *Stokes* | Cray 2 | 487.80 | 4 | 1.95 | 256 Mwords | 01/00/88 |
| *Piper* | CDC/ETA-10Q | 210.5 | 4 | 0.84 | 64 Mwords | 04/00/88 |
| *Reynolds* | Cray Y-MP | 317.46 | 8 | 2.54 | 32 Mwords | 08/30/88 |
| *Pierre** | Thinking Machines CM2 | 14 | 1,024 | 14.34 | 4 GB | 01/00/89 |
| *LaGrange* | Intel iPSC/860 | 60 | 128 | 7.68 | 4 GB | 01/00/90 |
| *Reynolds** | Cray Y-MP | 333.33 | 8 | 2.67 | 128 Mwords | 02/00/90 |
| *Pierre** | Thinking Machines CM2 | 14 | 1,024 | 14.34 | 4 GB | 02/00/91 |
| *Reynolds** | Cray Y-MP | 333.33 | 8 | 2.67 | 256 Mwords | 12/00/91 |
| *Boltzmann* | Thinking Machines CM5 | 128 | 128 | 16.38 | 32 MB per proc. | 01/00/93 |
| *Sigma* | Intel Paragon | 75 | 208 | 15.60 | 16 MB per proc. | 02/00/93 |
| *von Neumann* | Cray C90 | 960 | 16 | 15.36 | 256 Mwords | 03/00/93 |
| *Eagle* | Cray C90 | 960 | 8 | 7.68 | 256 Mwords | 05/00/93 |
| *von Neumann** | Cray C90 | 960 | 16 | 15.36 | 1,024 Mwords | 09/00/93 |
| *Babbage* | IBM SP-2 | 266 | 128 | 17.02 | 2 GB | 07/00/94 |
| *da Vinci* | SGI Power Challenge | ----- | 16 workstations | ----- | 512 MB | 09/00/94 |
| *Babbage** | IBM SP-2 | 266 | 160 | 42.56 | 256 GB | 10/00/94 |
| *da Vinci** | SGI Power Challenge XL | 360 | 32 | 11.52 | 512 MB | 08/00/95 |
| *Newton* | Cray J90 | 200 | 36 | 7.20 | 4 GB | 02/00/96 |
| *Piglet* | SGI Origin 2000/250 MHz | 500 | 8 | 4.00 | 2.048 GB | 03/17/97 |
| *Turing* | SGI Origin 2000/195 MHz | 390 | 24 | 9.36 | 7.168 GB | 03/28/97 |
| *Fermi* | SGI Origin 2000/195 MHz | 390 | 8 | 3.12 | 1.024 GB | 03/28/97 |
| *Hopper* | SGI Origin 2000/250 MHz | 500 | 64 | 32.00 | 16.38 GB | 12/31/97 |
| *Evelyn* | SGI Origin 2000/250 MHz | 500 | 8 | 4.00 | 2.048 GB | 12/31/97 |
| *Steger* | SGI Origin 2000/250 MHz | 500 | 256 | 128.00 | 65.536 GB | 03/27/98 |
| *Lomax* | SGI Origin 2800/300 MHz | 600 | 512 | 307.20 | 128 GB | 07/01/99 |
| *Lou* | SGI Origin 2000/250 MHz | 390 | 12 | 4.68 | 8 GB | 07/01/99 |
| *Ariel* | SGI Origin 2000/250 MHz | 500 | 8 | 4.00 | 4 GB | 05/01/00 |
| *Sebastian* | SGI Origin 2000/250 MHz | 500 | 8 | 4.00 | 4 GB | 05/01/00 |
| *Lomax** | SGI Origin 2800/400 MHz | 800 | 512 | 409.60 | 192 GB | 07/17/00 |
| Lomax II | SGI Origin 3800/400 MHz | 800 | 512 | 409.60 | 192 GB | 04/00/02 |
| *SN1 - 512* | SGI Origin 3000/400 MHz | 800 | 512 | 409.60 | 128 GB | 03/00/01 |
| *Bright* | Cray SV1e/500 MHz | 2000 | 32 | 64.00 | 32 GB | 05/25/01 |
| *Chapman* | SGI Origin 3800/400 MHz | 800 | 1,024 | 819.20 | 256 GB | 11/01/01 |
| *Chapman II** | SGI Origin 3800/600 MHz | 1200 | 1,024 | 1228.80 | 256 GB | 04/01/02 |
| **Data Assimilation Office** | | | | | | |
| *Jimpf0* | SGI Origin 2000/195 MHz | 390 | 64 | 24.96 | 16.38 GB | 06/26/97 |
| *Jimpf1* | SGI Origin 2000/195 MHz | 390 | 64 | 24.96 | 16.38 GB | 06/26/97 |
| *Helios1* | SGI Origin 2000/195 MHz | 390 | 8 | 3.12 | 2.048 GB | 03/28/97 |
| *Helios2* | SGI Origin 2000/195 MHz | 390 | 8 | 3.12 | 2.048 GB | 03/28/97 |
| *T1* | SGI Origin 2000/195 MHz | 390 | 8 | 3.12 | 2.048 GB | 03/28/97 |
| *T2* | SGI Origin 2000/195 MHz | 390 | 4 | 1.56 | 2.048 GB | 06/29/98 |
| *T3* | SGI Origin 2000/180 MHz | 260 | 4 | 1.04 | 1.024 GB | 06/24/97 |
| *T4* | SGI Origin 2000/180 MHz | 260 | 4 | 1.04 | 1.024 GB | 06/24/97 |
| *Lou* | SGI Origin 2000/195 MHz | 390 | 8 | 3.12 | 3.072 GB | 03/28/97 |
| *Gatun* | SGI Origin 2000/250MHz | 500 | 4 | 2.00 | 1.024 GB | 06/25/98 |
| *This* | SGI Origin 2000/250 MHz | 500 | 2 | 1.00 | 256 MB | 05/21/98 |
| *That* | SGI Origin 2000/250 MHz | 500 | 2 | 1.00 | 256 MB | 05/21/98 |
| *Kalnay* | SGI Origin 2000/250 MHz | 500 | 64 | 32.00 | 16.38 GB | 06/15/98 |
| *Dixon0* | SGI Origin 2000/250 MHz | 500 | 64 | 32.00 | 16.38 GB | 06/15/98 |
| *Sunrise* | SGI Origin 2000/300 MHz | 600 | 64 | 38.40 | 32.768 GB | 04/01/99 |
| *BJ* | SGI Origin 2000/300 MHz | 600 | 16 | 9.60 | 8.192 GB | 03/20/00 |

* upgrades

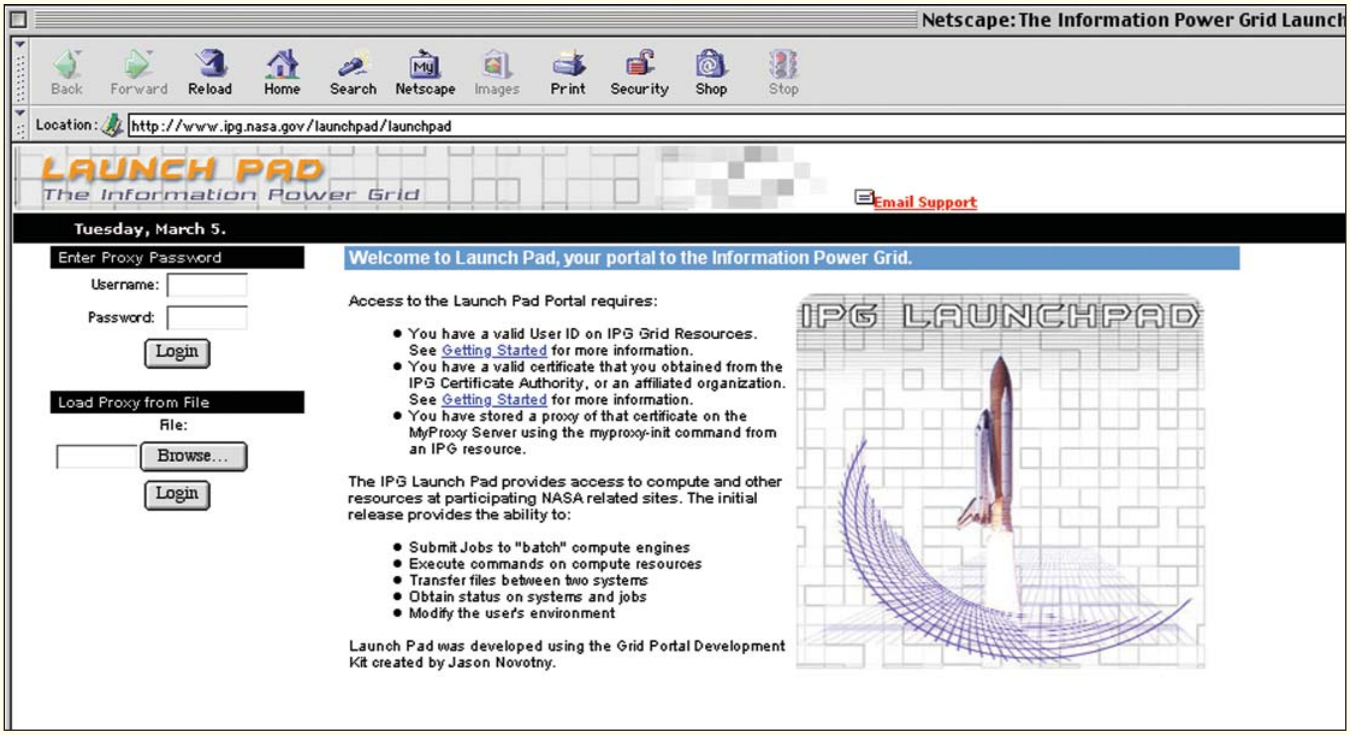


*LaunchPad's login screen can be accessed from any web browser, anywhere in the world. As long as users have a valid username and password, they have the freedom to launch jobs to any IPG resource. LaunchPad enables users to: submit jobs to "batch" compute engines, transfer files, view status on systems and jobs, execute commands on compute resources, and modify their user environment.*

*Continued from page 4*

"Our goal is to make LaunchPad modular and customizable so that researchers can build their own application-specific portal from it," says Myers. Each new resource in the grid environment adds a layer of complexity, and no matter what, security must always be addressed. While the web increases accessibility, it has the potential to make LaunchPad insecure.

## File Transfers You Can Trust

To ensure secure transmission of information, the web portal team established several requirements for logging into LaunchPad via a web browser. First, users must store their proxy certificate (an encrypted block of information about a user for authentication purposes) on a secure server (see *Gridpoints* Summer 2001, "Public Key Infrastructure: Get a Passport to Grid Country"). All certificates must be issued by an acceptable certificate authority (authorized people and software that sign certificates for authentication purposes), and researchers must have a valid user identification and a distinguished name (a unique combination of information that identifies an individual). User names must also be recognizable to the grid's database. In the near future, the team plans to expand on certificate options, alerting users when their certificate is about to expire, and giving users the option of selecting longer session times.

LaunchPad enables secure file transfer from any web browser, as long as the user is able to log in. The LaunchPad development team is now exploring options, such as the integration of web services to allow the maintenance of a consistent user interface, regardless of what is going on behind the scenes, and regardless of what changes in technology have taken place. "With web services, you could just execute a script that would access and move the files you needed, instead of having to move them directly," explains Myers. Web services is a new concept to the IPG community this year, and will be examined more extensively over the next nine months.

The next step for expanding the capabilities of LaunchPad is to focus on the development of application-specific portals. "If we can accomplish this, we will create an interface that is flexible enough to allow users to fully customize their view of LaunchPad," explains Myers. Once IPG LaunchPad permits researchers to access all resources on the grid, they will be able to accomplish anything they can currently do – all from the convenience of a web browser.

— *Holly A. Amundson*

For more information, visit the following websites:

NASA's Information Power Grid, visit:
http://www.ipg.nasa.gov

Grid Portal Development Kit:
http://dast.nlanr.net/Features/GridPortal

# Identifying Organic Molecules: The Link to Our Beginnings

**Computational chemists in the NAS Division are developing a theoretical model to generate accurate data of amino acids, which will help astronomers identify molecules observed in the interstellar medium.**

Astronomers study the sky night after night, searching for answers to questions about how life began on Earth and where it originated. Countless observations of organic molecules have been made by studying the interstellar medium. How will astronomers make sense of these observations? How can they trace their findings back to the beginnings of terrestrial life? Computational chemists Galina Chaban, Winifred Huo, and David Schwenke from the NASA Advanced Supercomputing (NAS) Division plan to compile a database to help astronomers identify the molecules they observe in space, and to reach accurate conclusions about their findings. Currently, little precise data on amino acids are available to the scientific community.

"Amino acids are fundamental building blocks of life. When astronomers try to find out if extraterrestrial life exists and how it is formed, a telltale sign will be the spectral signature of amino acids," explains Huo. "Amino acids and small peptides (strings of amino acids) have, in the past, been identified in meteorites, making the search for them in planetary environments especially relevant," adds Chaban. Since the Murchison meteorite crashed near Murchison, Australia in September 1969, scientists have identified within its composition, 92 different amino acids, including 19 of the 20 amino acids found in terrestrial biological systems. Discoveries such as this provide a strong lead to solving the mystery of life's origin.

## Breaking New Ground

To generate data for this catalog of amino acids, Chaban and her team will develop a theoretical model for infrared (IR) spectroscopy using both quantum mechanics and molecular mechanics, also known as force field calculations. This computational model will be continuously improved upon by comparing its result with experimental data. The experimental data are obtained using gas phase spectroscopy, which provides highly detailed output.

The team's work on IR spectroscopy was awarded monetary support through the NASA Ames Research Center Director's Discretionary Fund in summer 2001. For the duration of their funding period (through summer 2003), they plan to develop a theoretical model, and test it using accurate experimental results for small amino acids. Once the computational model can correctly produce data for small amino acids, it will be modified to accommodate larger amino acids. Research focus is on the vibrational spectra of amino acids and purines with high-intensity infrared radiation (in the

*The Murchison Meteorite crashed to Earth on September 28, 1969, near Murchison, Australia. The meteorite contains minerals, water, and complex organic molecules such as amino acids.*

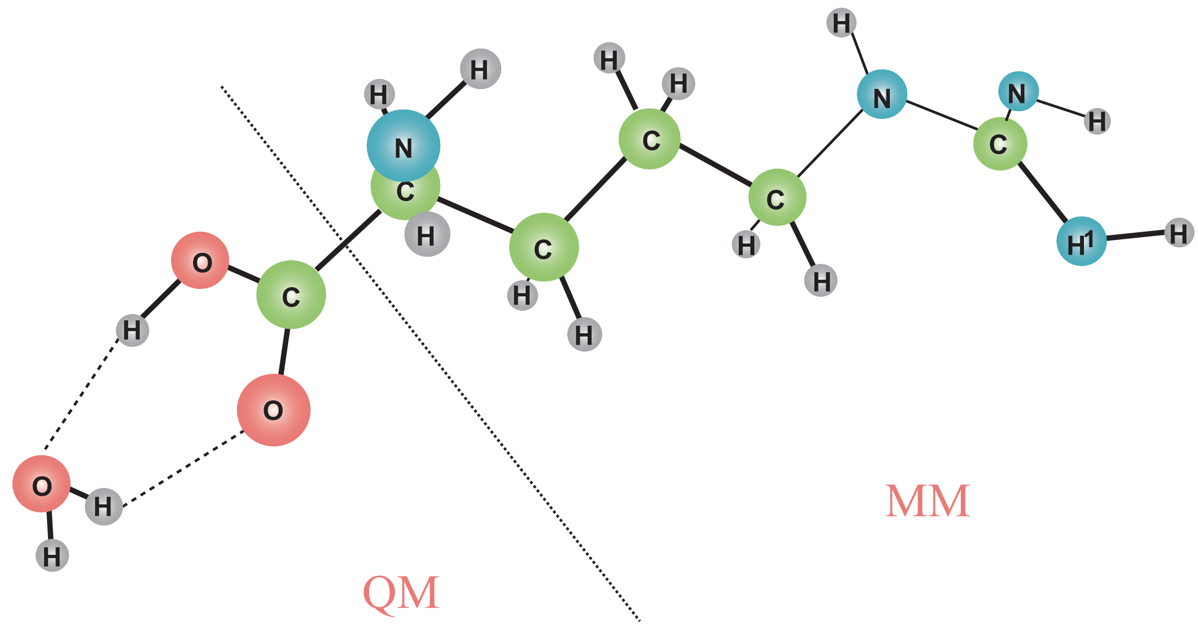


*Schematic illustration of the concept of the hybrid (Quantum Mechanics/ Molecular Mechanics or QM/MM) computational approach. Hydrogen bonds are not properly described by commonly used empirical force fields and need more accurate quantum-mechanical treatment. Thus, the region that includes hydrogen bonds is treated with* ab initio *(quantum mechanics), and the rest of the molecule (that contains standard chemical bonds) is treated by computationally less-demanding molecular mechanics potentials. The biggest challenge of this approach is proper treatment of borderline regions and cross-interactions between QM and MM atoms.*

region of 1,700 wave numbers, or $cm^{-1}$). These amino acids and peptides are of particular interest to astronomers because their presence in space would mean that important stages of prebiotic chemistry can take place in harsh interstellar conditions.

Traditionally, when studying amino acids through vibrational infrared spectroscopy (an important tool for identifying molecules), scientists use the molecular mechanics (MM) approach, which enables them to tackle large molecules (hundreds and thousands of atoms). With this method, however, accuracy is compromised. "The molecular mechanics method is like trying to study balls on a spring," explains Chaban. The quantum mechanics (QM) approach, on the other hand, is quite accurate, but only very small molecules (up to about 10 to 15 atoms) can be examined at one time. Since nearly all amino acids present in biological systems are rather large – 10 to 27 atoms (not to mention peptides, which are even larger than amino acids), a different approach must be taken – one that will provide both sound results, and have the ability to accommodate larger molecular systems.

Chaban, Schwenke, and Huo propose a unique vibrational IR method for studying amino acids, which combines both QM and MM (see figure of hybrid approach, above). "It's really complicated to combine two methods (MM and QM), but that's the only way to get accurate results with large molecules like amino acids and small peptides," says Chaban. This combined approach using hybrid potentials will enable the study of larger biological molecules (20-30 atoms), while also reducing the margin of error from 200-300 $cm^{-1}$, to 30-50 $cm^{-1}$. "The key part of the molecular system which has unusual interactions and requires more accurate treatment will be treated at a higher level (*ab initio*, or QM), while other parts of the molecule containing more standard chemical bonds will be treated with empirical (MM) potentials," explains Chaban.

The group is currently comparing the smallest amino acid, glycine (10 atoms), and a 13-atom glycine-water complex (see figures of glycine and glycine-water, page 8) using full *ab initio* (QM) calculations with the hybrid (QM/MM) technique. Their results have already increased in accuracy by one order of magnitude (an error of only 30-50 $cm^{-1}$). The team is aiming to increase the accuracy another order of magnitude – between one and three wave numbers. "The main challenge we are working on now is to develop a better treatment of the borderline regions of the molecule between the QM and MM regions," explains Chaban.

## The Experimental Side

Not only is it very difficult to obtain precise computational results for large molecules, but it is also extremely difficult to generate accurate experimental data. "Lack of experimental infrared data is due to the fact that such molecules are thermally unstable," explains Chaban. To sidestep the problem of thermal instability, Chaban proposes to collaborate with experimentalist Mattanjah De Vries, a professor at the University of California at Santa Barbara, who uses a laser desorption technique to bring molecules into the gas phase without undergoing decomposition.

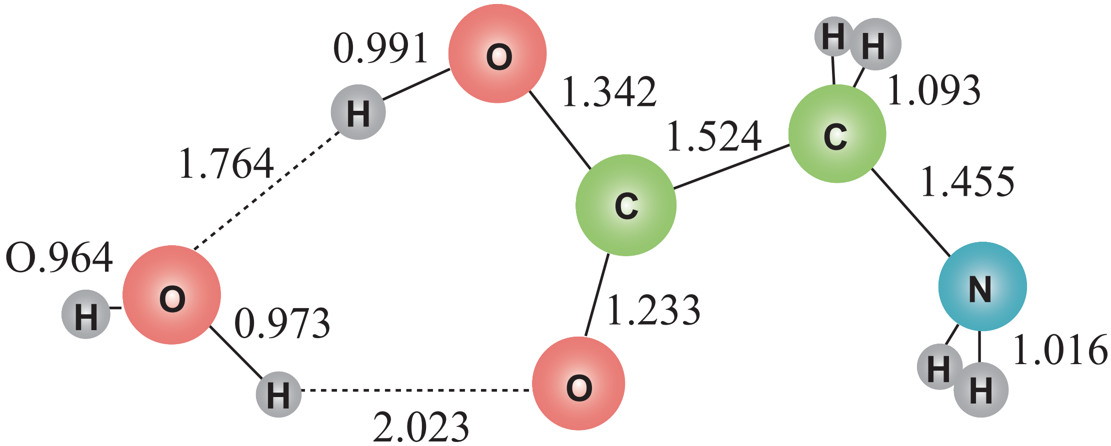


*Geometrical structure of the smallest amino acid, glycine, obtained using* ab initio *approach. Bond lengths are given in angstroms, bond angles in degrees. The structure corresponds to the lowest energy conformer of glycine. The spectroscopic predictions for glycine from the* ab initio *potential are in very good agreement with experimental results obtained in rare-gas matrix environments. This theory agrees even more closely with spectroscopic data for glycine in hydrogen droplets, where environmentl effects are much weaker.*

*Geometrical structure of the lowest energy conformer of the glycine-water complex obtained at the* ab initio *level of theory. The structure obtained using the hybrid QM/MM technique is very similar to the purely* ab initio *one. The intermolecular hydrogen bonding interactions between glycine and water are found to be very strong and to affect vibrational frequencies and infrared intensities of both the glycine and the water molecule in the complex to a very large extent. The computed vibrational spectra of complexes of amino acids and small peptides with water will allow scientists to guide experimental observations of these biological building blocks on icy surfaces of interstellar grains.*

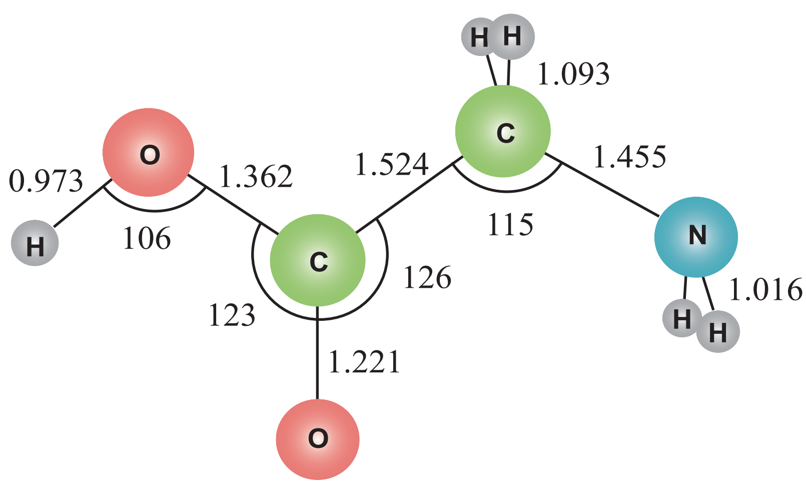


The theoretical model will be fine-tuned using experimental results as a control. Once the computations are comparable to experimental data, the model will be applied to a large group of molecules relevant to the study of the origin of life. The vibrational spectra of these biological molecules will then be used to build a database to aid astronomers. Initially, the database will hold information on 20 different amino acids – later, it will be expanded to hold records on approximately 90 amino acids and simple peptides.

Since the problems Chaban and her team are solving contain computations of thousands of energy points, parallelizing the code will enable faster execution of these calculations. A molecule with only 15 atoms, for example, involves calculations of tens of thousands of energy points. The team is currently using two computational chemistry software packages to run code on the 512-processor SGI Origin 2800 supercomputer, *Lomax.* The GAMESS electronic structure software package is used for *ab initio* calculations, and the TINKER molecular modeling package is used for molecular mechanics calculations.

"Mapping these compounds in space is of fundamental importance to understanding the origins, evolution, and distribution of life in the universe, as it can provide critical information for constructing models of chemical evolution under different astronomical conditions," says Chaban. Providing astronomers with a repository of data on amino acids will help to quickly identify molecules in the night sky, and begin tracing their chemical evolution – to the beginning of life on Earth.

*— Holly A. Amundson*

# Calendar of Events

### IEEE Symposium on Mass Storage Systems

**Adelphi, Maryland • April 15–18**

The 19th IEEE Symposium on Mass Storage will be held in cooperation with the 10th NASA Goddard Conference on Mass Storage Systems and Technologies at the Inn and Conference Center at the University of Maryland to provide a forum for discussion of issues relevant to the management of large volumes of data. For more information visit http://storageconference.org/2002

### International Parallel and Distributed Processing Symposium

**Ft. Lauderdale, Florida • April 15–19**

The International Parallel and Distributed Processing Symposium is designed for scientists and engineers from around the world to present their research and discoveries in all aspects of parallel computation. The symposium will include technical sessions, workshops, tutorials, and commercial exhibits. For conference information, visit: www.ipdps.org

### Grid Computing Planet 2002 Conference

**San Jose, California • June 17–18**

The Grid Computing Planet 2002 Conference and Expo will bring together grid technology researchers, developers, vendors and users to learn, exchange ideas and hear grid computing success stories. The event will feature an interactive exhibit floor enabling grid technology vendors to showcase their products. The conference emphasizes topics that will help attendees determine how grid technology can be utilized, what type of grid application will work best for them, and what tools are available to get started. Sessions will cover grid computing basics; related technologies such as distributed computing, P2P, clusters and caching; major grid initiatives; grid tools; grid security; grid computing and medical research; grid standards, policy requirements and implications; and grid case studies. More information is available at: www.intmediaevents.com/grid/spring02/

### Siggraph 2002

**San Antonio, Texas • July 21–26**

The 29th International Conference of Computer Graphics and Interactive Techniques is a forum for computer graphics scientists, artists, engineers, and educators. Technical presentations will include papers, panels, courses, applications, an educators program, web graphics program, and an applications lab. Conference information can be viewed at: www.siggraph.org/s2002

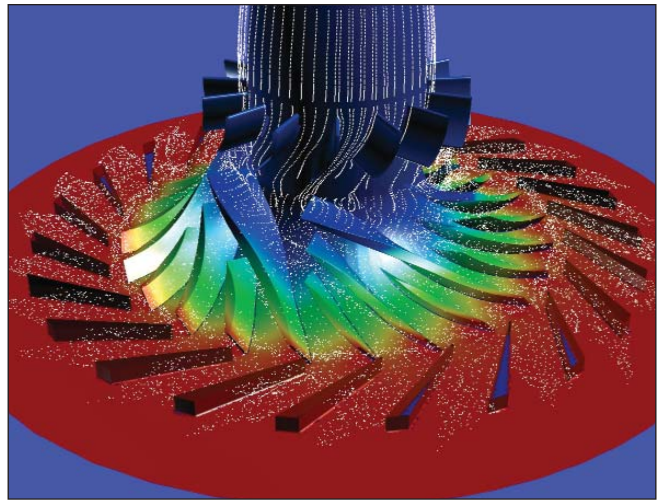

*Visualization of the entire fuel supply system of a liquid rocket engine, including high-fidelity unsteady turbopump flow analysis. The objective of this research is to provide a computational framework for design and analysis of the entire fuel supply system of a liquid rocket engine. This capability is needed to support the design of pump subsystems for advanced space transportation vehicles that are likely to involve liquid propulsion systems. (dataset by Cetin Kiris, visualization by David Ellsworth and Timothy A. Sandstrom)*

### The Grace Hopper Celebration of Women in Computing 2002

**Vancouver, B.C., Canada • October 9–12**

The Grace Hopper Celebration of Women in Computing 2002 is the fourth in a series of conferences designed to bring the research and career interests of women in computing to the forefront. Presenters are leaders in their respective fields, representing industrial, academic and government communities. Leading researchers present their current work, while special sessions focus on the role of women in today's technology fields. Visit www.gracehopper.org.

### SC2002

**Baltimore, Maryland • November 16–22**

"From Terabytes to Insights" is the theme for this year's supercomputing conference. SC2002 will bring together scientists, engineers, designers, and managers from all areas of high-performance networking and computing, showcasing the latest in systems, applications, and services. April 26 is the submission deadline for: Gordon Bell Awards nominations; panel participation; technical papers (extended abstracts); tutorials; and workshops. The conference website is: www.sc-2002.org

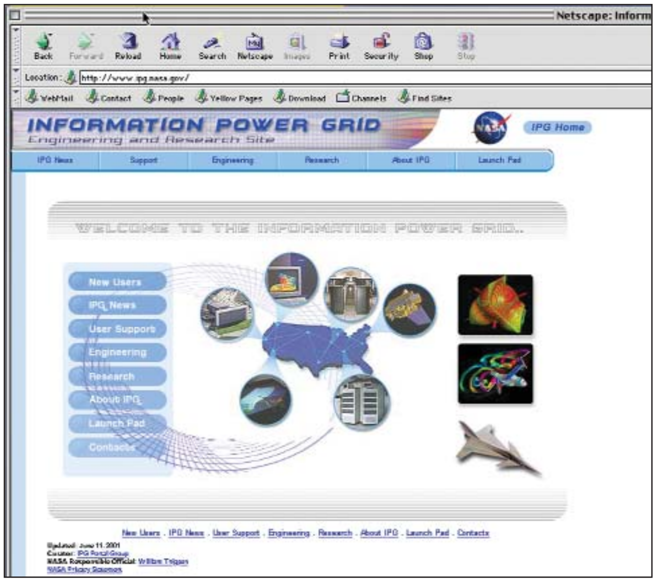


## Launching Jobs Into Grid Space

Information Power Grid users can now send jobs to any IPG resource through the recently developed grid portal, LaunchPad. See page 4.

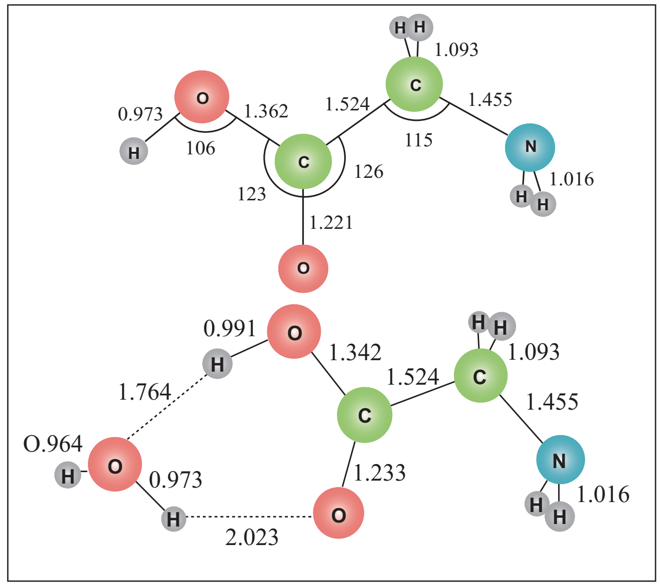


## Identifying Organic Molecules: The Link to Our Beginnings

Computational chemists in the NAS Division are developing a theoretical model to generate accurate data of amino acids, which will help astronomers identify molecules observed in the interstellar medium. See page 6.

**www.nas.nasa.gov/gridpoints**

National Aeronautics and Space Administration
**NASA Advanced Supercomputing Division**
Ames Research Center
MS IN258-6
Moffett Field, CA 94035-1000

Official Business
Penalty for Private Use $300

FIRST CLASS MAIL
POSTAGE & FEES PAID
NASA
Permit No. G-27

Address Correction
If the above address or mail stop is incorrect, please e-mail a current address to gridpoints@nas.nasa.gov, or write to:
Gridpoints Editor, MS 258-6, Moffett Field, CA 94035

PLEASE RECYCLE
Printed on recycled and recyclable paper with vegetable-based inks